National Aeronautics and Space Administration

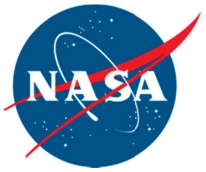

Single-Event Effect (SEE) Survey of Advanced Reconfigurable Field Programmable Gate Arrays

Gregory Allen
Jet Propulsion Laboratory
Pasadena, California

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

JPL Publication 11-18 12/11

National Aeronautics and Space Administration

Single-Event Effect (SEE) Survey of Advanced Reconfigurable Field Programmable Gate Arrays

NASA Electronic Parts and Packaging (NEPP) Program
Office of Safety and Mission Assurance

Gregory Allen
Jet Propulsion Laboratory
Pasadena, California

NASA WBS: 724297.40.49.11
JPL Project Number: 104593
Task Number: 40.49.03.03

Jet Propulsion Laboratory
4800 Oak Grove Drive
Pasadena, CA 91109

http://nepp.nasa.gov

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by the National Aeronautics and Space Administration Electronic Parts and Packaging (NEPP) Program.

TABLE OF CONTENTS

1.0 INTRODUCTION

This report is the result of funding by the NASA Electronic Parts and Packaging Program (NEPP) and, with respect to the testing performed on the XQR5VFX130 device, the combined efforts of members within the Xilinx Radiation Test Consortium (XRTC), sometimes known as the Xilinx Single-Event Effects (SEE) Test Consortium.

The NEPP Reconfigurable Field-Programmable Gate Array (FPGA) task has been charged to evaluate reconfigurable FPGA technologies for use in space. Under this task, the Xilinx single-event-immune, reconfigurable FPGA (SIRF) XQR5VFX130 device was evaluated for SEE. Additionally, the Altera Stratix-IV and SiliconBlue iCE65 were screened for single-event latchup (SEL).

In FY08, this task aided in the characterization of a 65 nm, 12-transistor static random access memory (SRAM)-based test chip developed by Xilinx. The test chip was designed to facilitate the acquisition and interpretation of SEE test data. In the following year, the Xilinx FX1, an engineering device that more closely resembles a commercial product was tested and reported on in [1]; this document is an augmentation of that report. The objective of FY10 SEE testing, reported on in [8], targeted the following high-level goals:

- Compare single-event upset (SEU) cross-sections for the XQR5VFX130 and FX-1 to cross-sections of the test chip (with focus on the fundamental building blocks of the device, including the configuration cells and block random access memory [BRAM]).
- Evaluate the device for any single-event functional interrupt (SEFI) and predict rates in the space environment.
- Begin to evaluate other structures, including the BRAM's error detection and correction (EDAC), digital signal processing (DSP) blocks, digital clock managers (DCMs), and multi-gigabit transceivers (MGTs).

The objectives of FY11 testing, the subject of this report, are as follows:

- Perform SEL screening of the final XQR5VFX130 silicon (previous SEL testing was performed on engineering samples)
- Complete the characterization of the flip-flops
- Continue characterization of IP blocks, specifically the DCMs and phase lock loops (PLLs)
- Evaluate other commercial reconfigurable devices for SEE

2.0 DEVICE OVERVIEW

2.1 65 nm SIRF Test Chip

The test chip was built on a 65 nm, triple-oxide United Microelectronics Corporation (UMC) process, fabricated on 2 µm epitaxial wafers. The die sizes were 4 × 6 mm and were bonded out to 208-pin grid packages. A special pedestal was introduced under the die in order to facilitate near glancing angle irradiations. The test chips contained 64 k, 12-transistor (12T) SRAM cells, single-event transient (SET) detectors and duration counters, and 12T flip flops (FFs) targeted for the configuration logic blocks (CLBs) (also used in the input/output (I/O) of the FX-1 devices).

2.2 65 nm FX-1 and XQR5VFX130

The FX-1 and XQR5VFX130 are 1.0 V, 65 nm, SRAM-based reconfigurable FPGAs. The devices are compatible with the commercially available Virtex-5 FX130T, and packaged in a 1752-pin ceramic, column grid array (CGA) package. Table 2.2-1 lists the architectural resources. The devices include the following programmable block types optimized for specific functions:

- Hardened-by-design (via dual-interlocked cells, or DICE) configurable logic blocks provide functional elements for combinatorial and synchronous logic, including configurable storage elements and cascadable arithmetic functions.
- The DSP slices provide advanced arithmetic and comparison functions, including multiply and accumulate. This block type is unhardened toward SEE.
- The block memory modules provide large 36-kbit storage elements of true dual-port RAM. The RAM is unhardened, but is mitigated with two-bit error detection and one-bit error correction.
- The clock tiles contain both DCM and PLL blocks that provide clock frequency synthesis and de-skew. These tiles are unhardened in FX-1.
- The MGTs provide high-speed serial transmission capability. These blocks are equivalent to the unhardened commercial MGTs.

Table 2.2-1. Architectural resources of the Xilinx XQR5VFX130 FPGAs.

	Description	Available Resources	Radiation-Hardened Implementation/Mitigation
CFG*	Configuration bits	49 x 10^6	12T cells (DICE)
BRAM	Block memory bits	10,985,472	EDAC
LOGIC	Slices (2, 6-input lookup tables/slice)	20,480	12T FF, 12T cells, SET filtering
DSP**	18 x 18 MACs	320	None
PPC***	PowerPC405 processors	0	None
CMT****	Clock managers	6	None
MGT	High-speed transceivers	18	None
IOBs	Input/output blocks	840	12T registers, TMR'ed Digitally Controlled Impedance Controller

* *Only real memory cells in the configuration bit stream are counted here (not counting BRAM)*

** *MAC = multiply-and-accumulate block for DSP*

*** *PPCs are not supported, but are still accessible in the FX-1; they were removed completely in the XQR5VFX130*

**** *Clock management tiles contain two DCMs and one PLL*

2.3 Altera Stratix-IV

The Altera Stratix-IV is a 0.95 V, 40-nm device built on the Taiwan Semiconductor Manufacturing Company (TSMC) process technology. The Stratix-IV devices are available as four family variants: the GT and GX, GS, and E supporting high bandwidth, DSP, and high logic application devices, respectively [2]. The Altera EP4SGX230KF40 FPGA was specifically targeted for this testing. The EP4SGX230KF40 includes the following programmable block types optimized for specific functions: 747 user I/Os, 8 PLLs, 1288 embedded 18x18 multiplier blocks, 91,200 adaptive logic modules, 228,000 logic elements, and 16 high-speed transceivers. The devices were commercial engineering devices in a F1517 ball grid package. Similar to the Xilinx devices, these flip-chip devices were thinned in order to achieve ion penetration.

2.4 SiliconBlue iCE65

The SiliconBlue iCE65 and Altera Stratix-IV devices are commercially available, reconfigurable FPGAs. The SiliconBlue iCE65 family of FPGAs is fabricated on a 65-nm complementary metal-oxide semiconductor (CMOS) process with an internal core voltage of 1.0 or 1.2 V, depending on speed grade. They are relatively low-power devices, typically drawing current in the micro-amp to milliamp range, and are capable of an internal performance of up to 256 MHz [9]. The specific device targeted for this testing was the iCE65L04VF-LCB284C, D/C 0852. These fully reconfigurable, low-power devices are best suited for small microcontrollers or glue-logic implementation. The tested SiliconBlue devices were packaged face-up in a plastic BC284 package. The devices were etched using an acid-etching machine in order to expose the die and achieve ion penetration. The iCE65L04VF tested has 3,520 logic cells, 200K system gate equivalent, 80K embedded RAM bits, and a total of 533K configuration bits.

3.0 FACILITIES, EXPERIMENTAL SETUP, AND DEVICE PREPARATION

3.1 Facilities

Heavy ion SEE measurements were performed at the Radiation Effects Facility located at Texas A&M's Cyclotron Institute (TAM) and at the Berkeley Accelerator Space Effects as part of the 88-inch cyclotron at Lawrence Berkeley National Laboratory (LBNL). Each facility uses an 88-inch cyclotron to provide a range of ion beams and energies.

3.2 Experimental Setups

Figure 3.2-1 shows the heavy ion test setup in vacuum. Vacuum irradiations were performed for all tests at LBNL, for 15 MeV/amu Au irradiations at TAM, and for low-energy irradiations at Crocker Nuclear Laboratory (CNL). Details of the test board used for all experiments are found in [3]. When in vacuum, five 40-pin bulkheads were used to run five of the six communication cables through the vacuum chamber. The sixth was run through the 50-pin D-Sub connector provided by TAM. Three parallel cables were also sent through the 50-pin connectors (one for a device under test [DUT] readback Parallel-IV cable, one for a motherboard/DUT design programming Parallel-IV cable, and one for the temperature monitoring circuit used during SEL testing). A mounting platform with integrated power breakout cables was used for mounting the motherboard to the rotating chassis in the vacuum chamber, and for extracting the four power supplies from the 40-pin cable. The four supplies were sent through the vacuum chamber bulkhead over Bayonet Neill Concelman (BNC) connectors, and then re-integrated to the 40-pin cable. Force and sense were tied together at the power supply (HP6629) for all four supplies, and provided the necessary 2.5 V, 3.3 V, and 3.3 V I/O for the motherboard; the last supply was used to control heater strips attached to the back of the daughter card (when performing SEL testing). The receiver/driver cards were powered by the 3.3 V of the motherboard I/O. The 5 V for the Parallel-IV cables and temperature sensor circuit were powered by an external Agilent E3610A, and also run through a BNC bulkhead. Typically, the DUT power supply was an HP6623 or HP6624, which are capable of providing supply currents in the 5 and 10 Amp range. Supply one provided 2.5 V to the auxiliary voltage supply (V_{AUX}) while supply two provided 1.0 V to the internal voltage (V_{INT}), and supply three provided 3.3 V to two of the I/O voltage (V_{CCO}) DUT I/O banks that talk to the motherboard. Force and sense were tied together at the bulkheads on supplies one and three (which were run through BNC bulkhead feedthroughs). High current cables were used to run force for supply two in through a 40-pin cable and bulkhead connector, and were separated back into banana cables with a second custom 40-pin connected to an additional banana cable (20 pins were used for power and 20 for ground). Sense for supply two was sent through a BNC over banana cables and connected to force at the daughter card.

The setup for in-air testing was essentially the same as in vacuum, the main exception being that the adapted connections for getting through the bulkheads were discarded. Also, USB programming cables were used via high-speed hubs for the in-air irradiations.

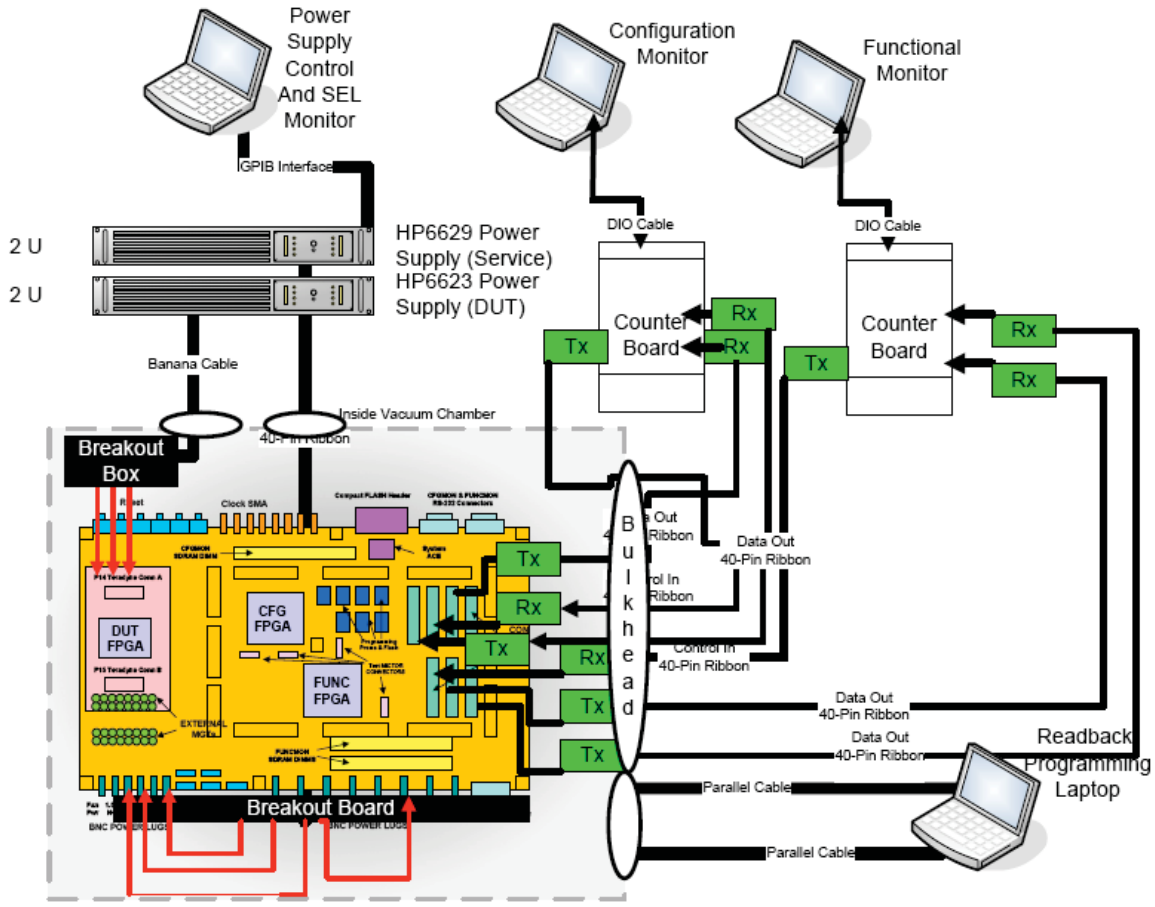

Figure 3.2-1. Experimental setup in vacuum at Texas A&M. Everything outside of the vacuum chamber border is set up in the SEE user room.

Evaluation boards were used for both the Altera and SiliconBlue testing. In the case of the Altera devices, a single evaluation board was procured and the FPGA was replaced with a socket. The regulated core voltage was bypassed and replaced with an adapter to supply the voltage with a HP6624 power supply. Commercially available software that provided off-chip (i.e., the control circuitry was not in the beam) temperature and power monitoring was used in conjunction with custom Visual Basic software described in Section 4. Three separate evaluation boards were used for the SiliconBlue testing; all supply voltages were supplied via an HP6629 power supply. Figure 3.2-2 shows the SiliconBlue (top) and Altera (bottom) evaluation boards.

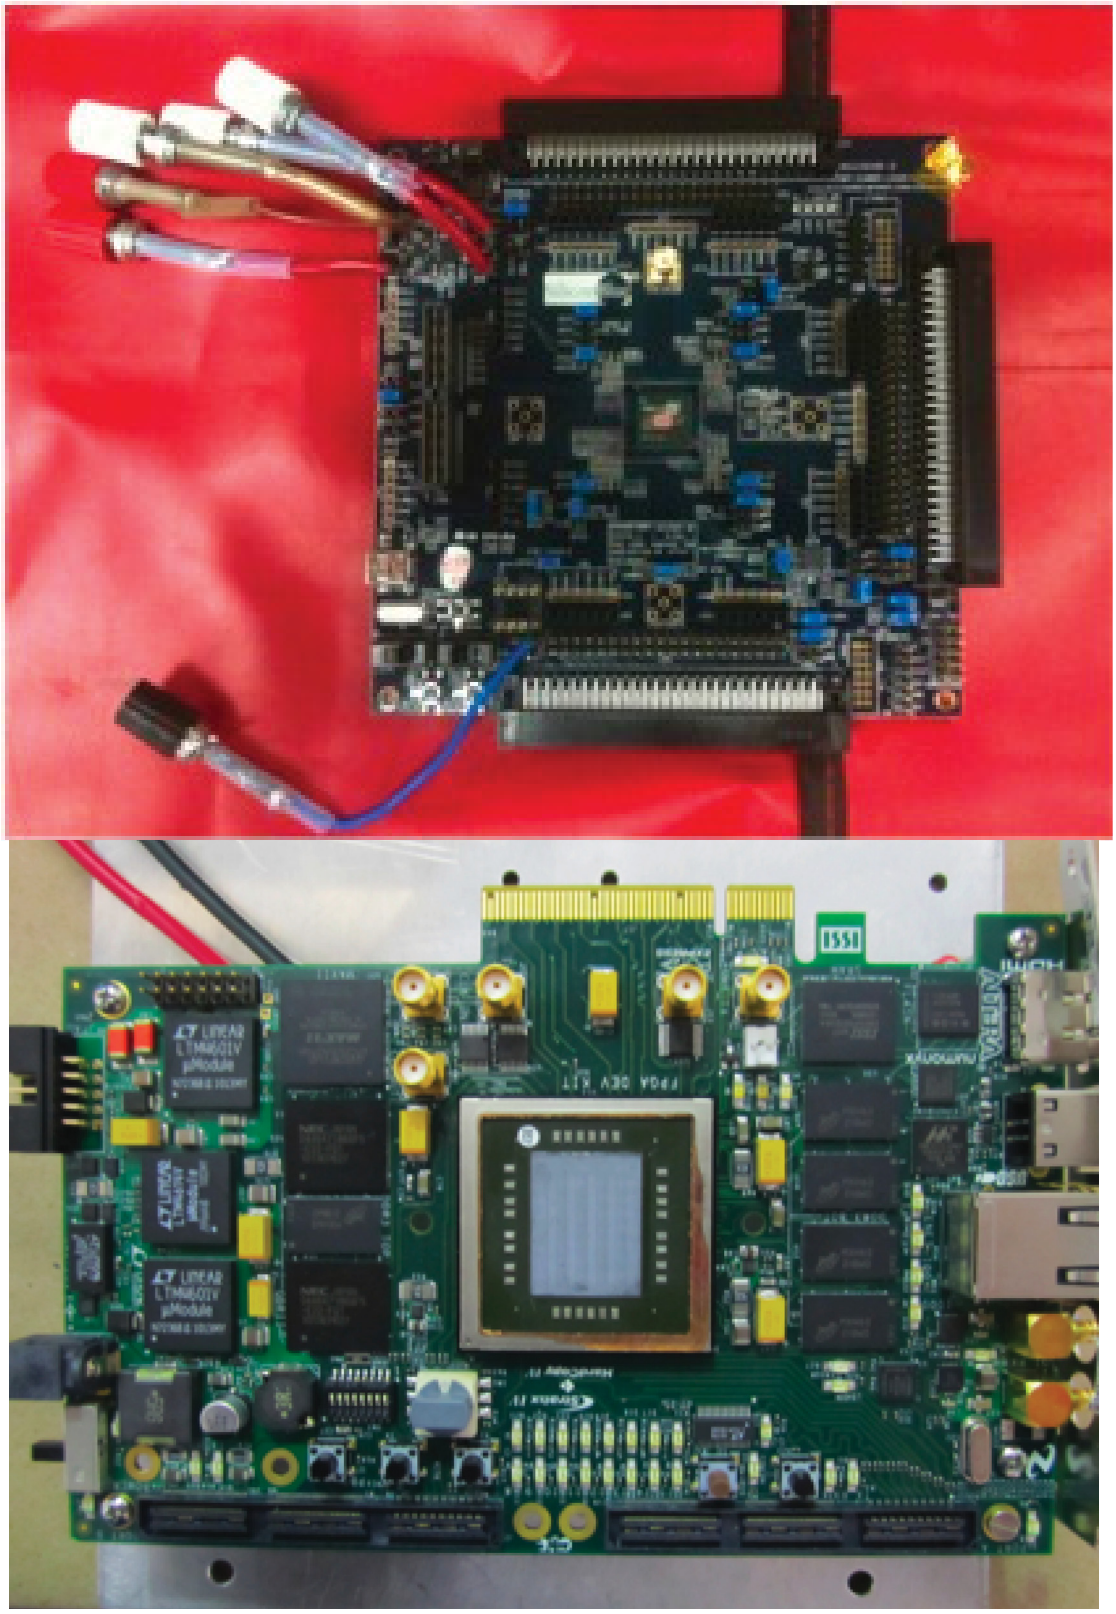

Figure 3.2-2. The SiliconBlue (top) and Altera (bottom) evaluation boards.

3.3 Device Preparation

Because the XQR5VFX130s and Stratix-IV are only offered in flip-chip packaging, irradiation is done through the backside of the silicon substrate, which is exposed after delidding. In order to reach the active layer with a high-linear-energy-transfer (LET), short-range heavy ion, the backside of the silicon must be thinned to approximately 100 µm or less. This is done with a high-precision, micro-milling machine. The effective range measurement reported in this document is the residual silicon-equivalent penetration depth after the ion penetrates the thinned backside and the epitaxial layer, i.e., the residual effective range after exiting the active layer; see Figure 3.3-1.

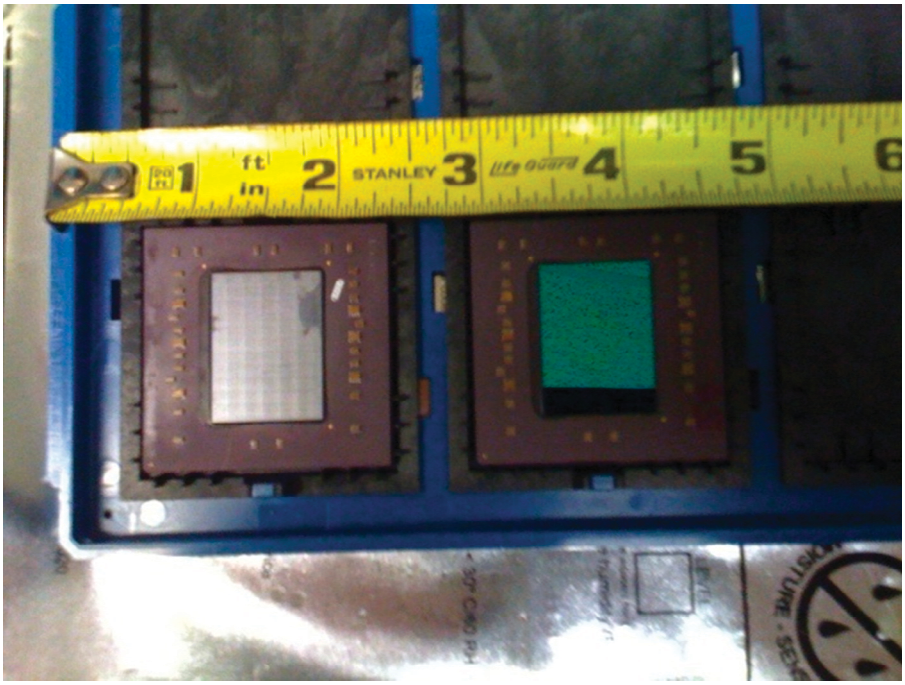

Figure 3.3-1. Two FX-1 devices; the device on the left has been thinned to approximately 100 µm of remaining silicon. The device on the right is un-thinned.

4.0 SINGLE-EVENT EFFECTS RESULTS

4.1 Xilinx 65 nm Test Chip SEE Overview

This document does not provide in-depth analysis of the 65 nm test data. Full analysis of previous tests is available in [4], and can be used to calculate upset rates for the configuration cells. This document instead provides an overview of the test approach and data analysis to provide the reader with insight into characterizing the Xilinx XQR5VFX130 device regarding the configuration cells. Typically, we characterize SEE susceptibility of digital devices from the standpoint of the single-node, or single sensitive volume. A device, such as a memory, contains a set of sensitive volumes; an error occurs if any one or more of these sensitive volumes collects enough charge to upset the corresponding bit. The concept of single-node refers to the model in which the upset of a given bit depends only on the conditions seen by the single sensitive volume, independent of other sensitive volumes. In a 12T, SEE-hardened SRAM cell (a DICE cell), SEE tolerance has been built in such that a bit upset will require a pair of sensitive volumes to be upset simultaneously, meaning a sufficient amount of charge is collected in each sensitive volume from the same particle. This document refers to this as the dual-node problem.

A dual-node device exhibits complex directional dependencies when measuring the cross-sections with a particle accelerator that are not seen with single-node devices. A multi-node cross-section becomes a function of LET, DUT tilt (the beam angle in reference to the DUT surface-normal axis), and DUT rotation (the DUT's azimuth angle). One can expect a negligible upset cross-section at normal incident beam if the special separation of the two sensitive nodes is larger than the typical drift and diffusion length in the device. At normal incidence, only the rare event in which a nuclear reaction product with enough charge deposition and appropriate directionality will cause an upset. As the tilt increases toward 90º, and the rotation oriented such that the dual-nodes begin to line up with the particle path, the probability that a critical charge will be deposited in each of the sensitive volumes increases. Plots must then be produced as seen in Figure 4.1-1 for various tilts.

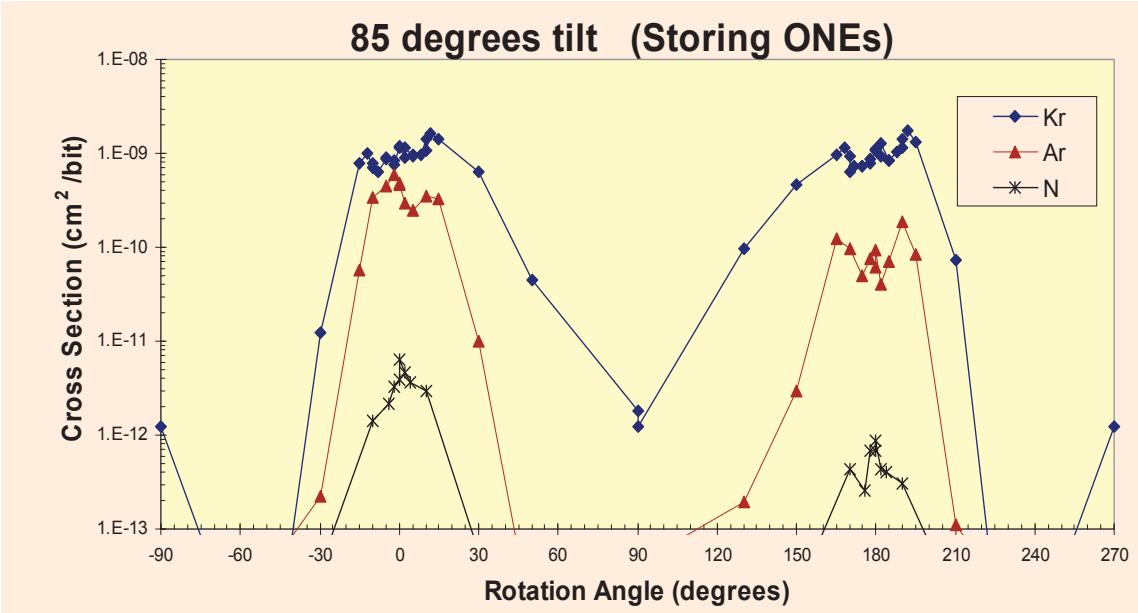

Figure 4.1-1. Various LET, per-bit cross-sections versus rotation angle for a single tilt. A plot such as this is repeated for several tilt angles.

Two 65 nm test chips were developed. Data collected from configuration cells on each chip showed similar SEE responses. The first test chip contained several SRAM structures with variable sensitive volume spacing, in order to determine an appropriate trade-off between speed and SEU hardness. Once an acceptable geometry was determined, the test chip described in Section 2.1 was developed to acquire cross-section data to estimate an upset rate for the configuration cells. The test chip was required due to the flip-chip nature of the FX-1 FPGA. The near-glancing angles at which the SRAM cells are most sensitive could not be acquired with a flip-chip device due to ion range issues for increasing LET.

A new model and methodology had to be developed in order to calculate space rates, as the traditional rectangular parallel-piped (RPP) model cannot be used for dual-node devices. Two error rate calculation methodologies were developed and compared. A Monte Carlo simulation technique and an analytical model were used to calculate per-bit error rates of configuration cells on the test chip. Each methodology came up with similar rates for a galactic cosmic ray (GCR), solar minimum environment: 2.1×10^{-10} per bit-day for one-to-zero bit flips and 5.3×10^{-11} per bit-day for zero-to-one bit flips.

4.2 Xilinx XQR5VFX130, Altera Stratix-IV, and SiliconBlue iCE65 SEL Results

Full SEL characterization of the production mil/aero Xilinx 5QV and commercial Altera and SiliconBlue devices took place at TAM in July 2011. In general, SEL characterization was performed in accordance with ASTM guideline F 1192-00 [5]; no latchup events were recorded for any of the three devices. Table 4.2-1 shows the SEL run parameters. The DUTs were heated to an elevated die temperature of approximately 125°C for the Xilinx device and 85°C for the Altera devices, as determined by internal temperature sensing diodes and custom Visual Basic software. In all cases, the DUTs were biased with specified maximum voltages.

For the purpose of these experiments, the accepted definition of a latchup was any sudden high current mode resulting from the test run that required a power cycle of the DUT in order to recover. During the test runs, the DUT core, I/O voltages, and dynamic current consumption were captured and recorded in a running log (strip chart). Maximum current triggers were set on the power supplies in the event of a latchup condition that would result in excessive current draw. Due to the high fluxes and total fluences used for the latchup testing, it was expected that the DUT would lose its programming early in the run and would likely be subject to multiple SEFI conditions during the run. The purpose of the experiment was to demonstrate hardware survivability and soft recovery without the need for a device power cycle. Therefore, the test procedure adopted was as follows:

1. Program and readback to verify DUT configuration memory.
2. Heat DUT to an elevated temperature.
3. Record initial temperature, voltage, and current conditions.
4. Irradiate the DUT to at least 10^7 particles/cm^2.
5. Record DUT power and temperature conditions during irradiation run.
6. Program and readback to verify DUT configuration memory after end of irradiation (Xilinx devices only).
7. Perform functional verification.

Because the bottom of the silicon is solder “bumped” to a fully populated ball-grid package, it is difficult to heat the device enough for latchup testing with an external heating element. In order to obtain the target temperature, the devices were configured with a “heater” design meant to increase dynamic current consumption sufficient to heat the transistor junctions to a desired temperature.

The heater design is a long shift-register chain of CLB FFs. Typically, this chain is long enough to consume more than 75% of the available device resources. A one-bit counter feeds the start of the register chain so that alternating ones and zeros advance through the chain with each clock pulse. In order to obtain a high enough frequency to meet the dynamic consumption requirements, a DCM is used to multiply the input clock frequency. The typical internal clock frequency was ~80 MHz.

In the case of the SiliconBlue devices, Californium-252 fission fragments were used to screen for SEL on delidded devices in vacuum. While the devices were biased at the maximum specified voltages, the devices were not heated beyond the operational temperature in vacuum. The face-up test devices were de-packaged using an acid-etching machine, and banana plugs were added to the evaluation boards in such a way as to bypass the board's voltage regulators and supply power directly to the DUT.

The dynamic current draw was less than the measurable threshold of the power supply, 100 μA, and the SEL threshold was set to 50 mA. While no SELs were observed over a two-day exposure, the current increased linearly, at an approximate rate of 3.5 mA/hr. The linear increase in current consumption indicates contention in the device due to SEU in the configuration memory. A total of three devices were tested with Xenon ions to a total fluence of 1×10^7 ions/cm^2 each. During heavy ion testing, a similar current ramping was observed as seen in the Californium testing; however, no SEL was observed. These devices were tested at the specified maximum voltages and at room temperature. Table 4.2-1 summarizes the SEL runs.

As a reference for the reader, the previously tested Altera Stratix-II [6] exhibited severe SEL sensitivity. Previous work shows that due to scaling effects, SEL may no longer be an issue for these devices [7]. A total of three devices were tested with Xenon ions to a total fluence of 1×10^7 ions/cm^2 each. These devices were tested at the specified maximum voltages and at elevated temperature, typically near 85ºC.

Table 4.2-1. Texas A&M 15 MeV/amu Latchup Test Data

Device	ION	LET MeV-cm^2/mg	Effective LET MeV-cm^2/mg	Effective Range μm	Fluence ions/cm^2	SEL
			XQR5VFX130			
SN 515	Au	93.5	145.5	34.6	2.00x10^7	0
SN 514	Au	93.5	145.5	34.6	2.00x10^7	0
SN 592	Au	93.5	145.5	34.6	2.00x10^7	0
			iCE65L04V			
SN 1	Xe	52	83	57.8	1.00x10^7	0
SN 2	Xe	56	83	42.8	1.00x10^7	0
SN 3	Xe	56	83	42.8	1.18x10^7	0
			EP4SGX230KF40			
SN 559	Xe	41.5	104.1	113	1.00x10^7	0
SN 558	Xe	41.5	112	113	1.00x10^7	0
SN 557	Xe	41.5	112	113	1.00x10^7	0

4.3 XQR5VFX130 Configuration Bit SEU

Due to the angular dependence of the configuration cells discussed in Section 4.1, and the flip-chip geometry of the XQR5VFX130, full configuration cell characterization cannot be performed on this device. However, normal-incident and limited low LET (limited by ion range) cross-section data can be measured and compared to the test chip data. At shallower angles and using higher energy ions and/or lighter ions, it is possible to obtain upset data on the FPGA itself to compare to the test chip data. Recent testing was performed at the TAM cyclotron to obtain such data. Unfortunately, comparable test chip data for that low an LET and with statistical significance has not been acquired.

While a large majority of the 49 million configuration cells follow the SEU response of those measured in the test chip, there are a few discrepancies. There are two types of SEE responses observed in the configuration memory beyond the typical 12T configuration cell upset. The two cell types are dubbed *INIT/CAPTURE bits (ICb's)* and *Weakly Loaded Common Address Line bits (WLCALb's)*. Using upset-hardened-by-design techniques, it is possible to drive the direct upset rate down to such a low level that SET-triggered upsets dominate. In the static case, SETs coincident with a clock edge are ruled out, but SETs on asynchronous control lines (say reset or write signals) or on the clock lines themselves can cause an upset. ICb's have no impact on the functionality of the design; they can simply be masked out by the FPGA's configuration management and ignored. It should be noted that if a readback-scrub configuration management scheme is implemented, "false positives" will occur if proper masking is not used. Conversely, WLCALb's can affect the functionality of a design. Sets of configuration cells share common address and control lines. An SET on a control line for the right condition may cause configuration cells to have the arbitrary value on the data lines loaded into the cell. Several sets of data were taken at TAM in July 2011, with a mask that eliminated ICb's from the device readback. The goal was to extract cross-sections WLCALb's. The data was mostly taken at normal incidence, in order to eliminate SEU from direct ionization of the dual node cells. Preliminary analysis has been performed and a conservative rate of approximately one configuration bit per device per two years (for the standard GCR solar min, behind 100 mils Al) was calculated.

4.4 Xilinx XQR5VFX130 Single-Event Transient Testing

SET testing took place on the FX-1 and production XQR5VFX130 at TAM and LBNL over several test campaigns during FY10 and FY11. The cognitive engineer for this testing was George Madias of Boeing; Eric Miller, also of Boeing, performed a great deal of the analysis. The test consisted of a simple shift register that was clocked at various speeds and with different SET filter settings (0 ps, 400 ps, or 800 ps), different clock speeds (1.5 MHz, 100 MHz, and 200 MHz), and varying levels of logic.

Three pattern modes were used: pattern 0—checkerboard pattern, pattern 1—all ones, and pattern 2—all zeros. The all-zeros pattern is only sensitive to data path hits, i.e., clock and reset hits will not affect the output. The all-ones pattern is sensitive to data path transients and clear transients. Finally, the checkerboard pattern is sensitive to clear, clock, and data transients. Several different types of combinatorial logic designs were implemented to see if there was an effect due to varying levels of combinatorial depth, fan-in, and look-up table (LUT) implementation (dubbed sensitive and non-sensitive, which refers to whether or not SET on unused LUT inputs effect the output). Figure 4.4-1 exemplifies the frequency dependence observed on the data path with four levels of combinatorial logic. Figure 4.4-2 exemplifies how clock transients dominate the upset mode at lower frequencies up to 200 MHz where the data path and clock transients become comparable. Figure 4.4-3 shows that increasing the LUTs increases the cross-section linearly. The all-ones pattern, serial combinatorial logic implementation was used at 200 Mhz. Figure 4.4-4 exemplifies the relative effectiveness of the SET filter settings, reducing the cross-section by about an order of magnitude. Note, however, that the filters do not change the LET threshold, which is somewhat surprising. Finally, Figure 4.4-5 shows the relative effectiveness of the SET filters for the varying types of LUT configurations. Note that the SET filters are less effective for the parallel sensitive configuration.

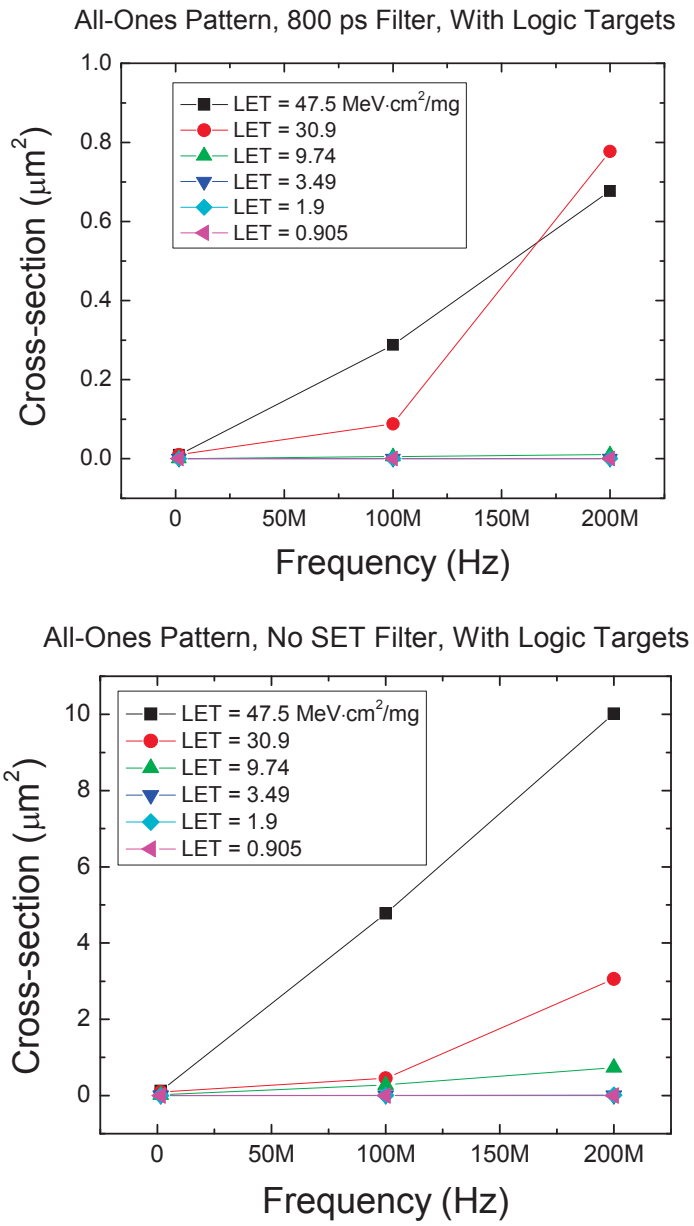

Figure 4.4-1. Shows the frequency dependence of the flip flops with (top) and without (bottom) SET filters.

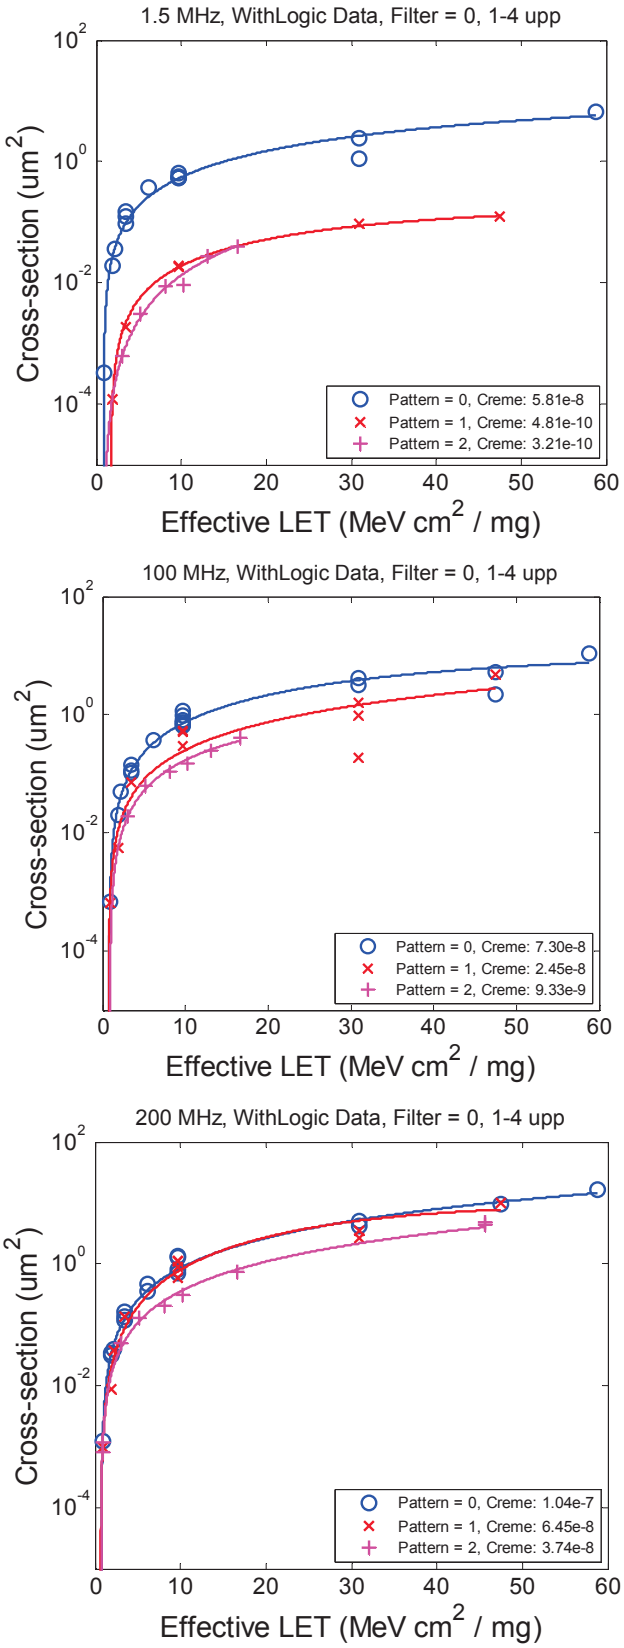

Figure 4.4-2. Shows the frequency dependence of pattern 0 (checkerboard), pattern 1 (all ones), and pattern 2 (all zeros) for 1.5 MHz (top), 100 Mhz (middle), and 200 MHz (bottom)

LET (MeV.cm2/ mg)	Cross-section (um2)		
	No Logic	2 LUTs	4 LUTs
0.905		0.0005	0.0010
1.9	0.0021	0.0048	0.0088
2.19	0.0042		0.0383
9.74	0.5532		1.1173
30.9	1.4151		3.0628
47.5		6.5091	10.0193

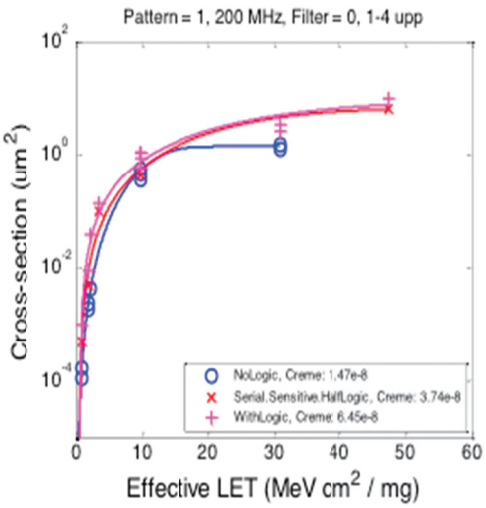

Figure 4.4-3. Shows the effect of various levels of combinatorial logic. Pattern 1 (all ones) is shown here at 200 MHz.

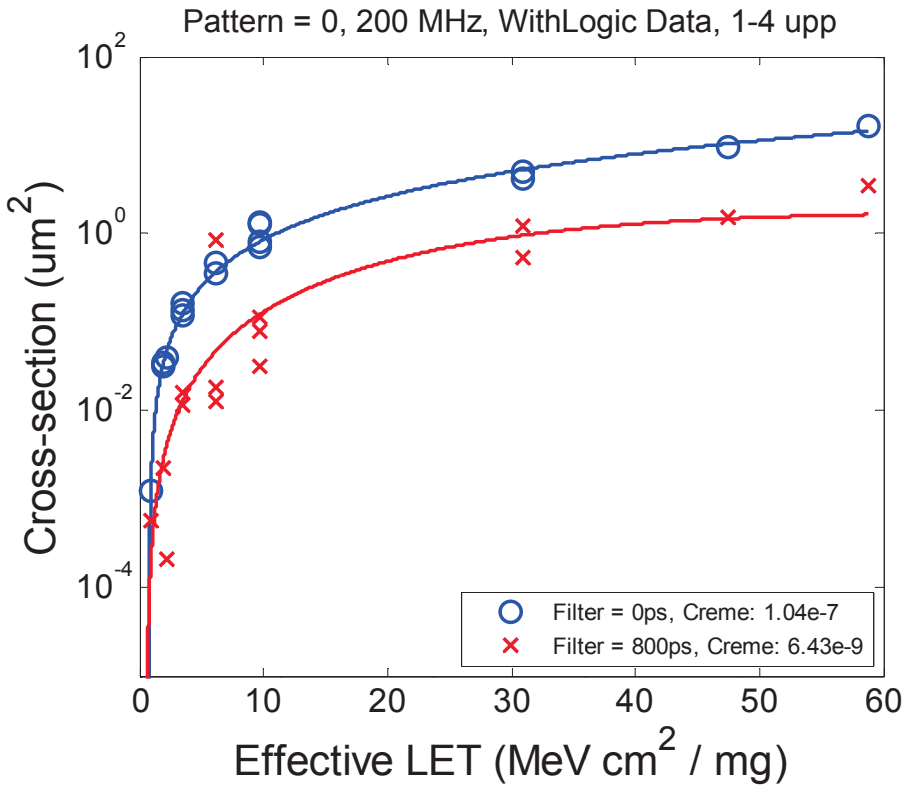

Figure 4.4-4. Shows the effectiveness of the SET filter for the four levels of serial combinatorial logic case running at 200 MHz.

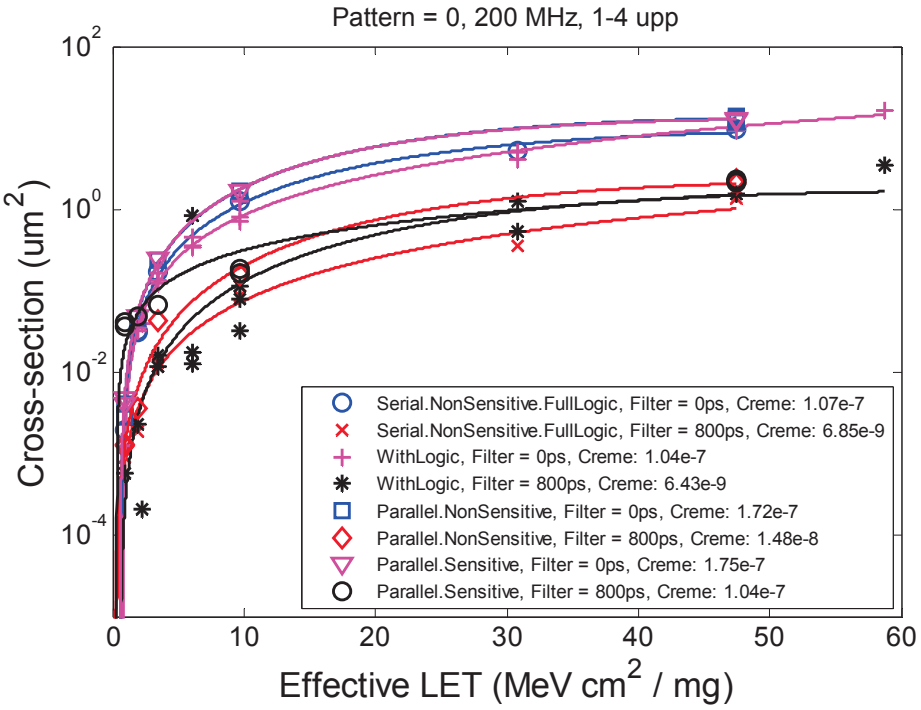

Figure 4.4-5. Shows cross-sections for shift registers with and without filters, and for varying types of combinatorial logic (serial and parallel).

4.5 XQR5VFX130 Clock Management Tiles

This section focuses on the results of the clock management tiles (CMTs) of the Xilinx XQR5VFX130. Each CMT consists of one PLL and two DCMs. Preliminary testing of the CMTs took place at TAM. The testing was two-fold. The first set of tests simply consisted of instantiating either a PLL or DCM, and monitoring the output with a watchdog timer as well as the lock pins. When the test engineer observed a functional failure, either the lock pin was de-asserted or the clock output was arrested, then the beam was paused and the recovery mechanism was investigated. A second test consisted of the DCM arrangement shown in Figure 4.5-1.

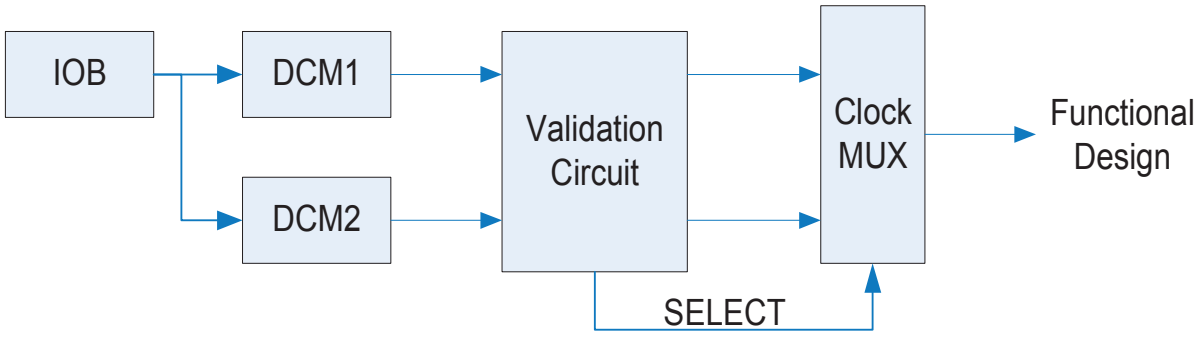

Figure 4.5-1. Top-level block diagram of the DCM/PLL characterization circuit for the Xilinx XQR5VFX130. This implementation was instantiated twice; each instantiation is referred to as a separate bank.

The validation circuit compared the DCM output of DCM1 to an expected value, and counted the clock cycles if that value was incorrect. If a clock output was deemed incorrect, DCM2 would take over as the primary clock output until DCM1 recovered. This implementation was instantiated in two separate I/O banks; the input clock was 33 MHz. The lock signal was brought out and monitored, as well as each instance of the CLK0 output.

For both the DCM and PLL (subsequently to be referred to generically as Clock), the observable error signatures consisted of altered frequency, clock glitches, or completely arrested functionality. With respect to the arrested functionality, the recovery mechanism consisted of only a reset to the CMT and scrub. Figures 4.5-2 through 4.5-6 summarize the results of the testing. It should be noted that the characterization was performed with the GLUT mask disabled, meaning dynamic reconfiguration (DRP) bits were being scrubbed. A multiple-bank switch error is defined as switching between Clocks multiple times (greater than three) in a given instantiation. This is most likely due to an upset in the output clock path. Figure 4.5-2 shows these data. A second type of error is the both-bank multiple switch event. This is similar to the multiple-bank switch event, except that it occurs in both instantiations at the same time. This is most likely due to an upset in the input clock path. Figure 4.5-3 shows these data. A unique single switch event is a switch between Clocks that happens less than three cycles, and occurs in only one of the instantiations. Figure 4.5-4 shows these data. A single bank switch is similar to a unique single switch event, except that it can occur in one or both of the instantiations. Figure 4.5-5 shows these data. Finally, there is the loss-of-lock (LOL) event, where the lock output signal is de-asserted, as shown in Figure 4.5-6.

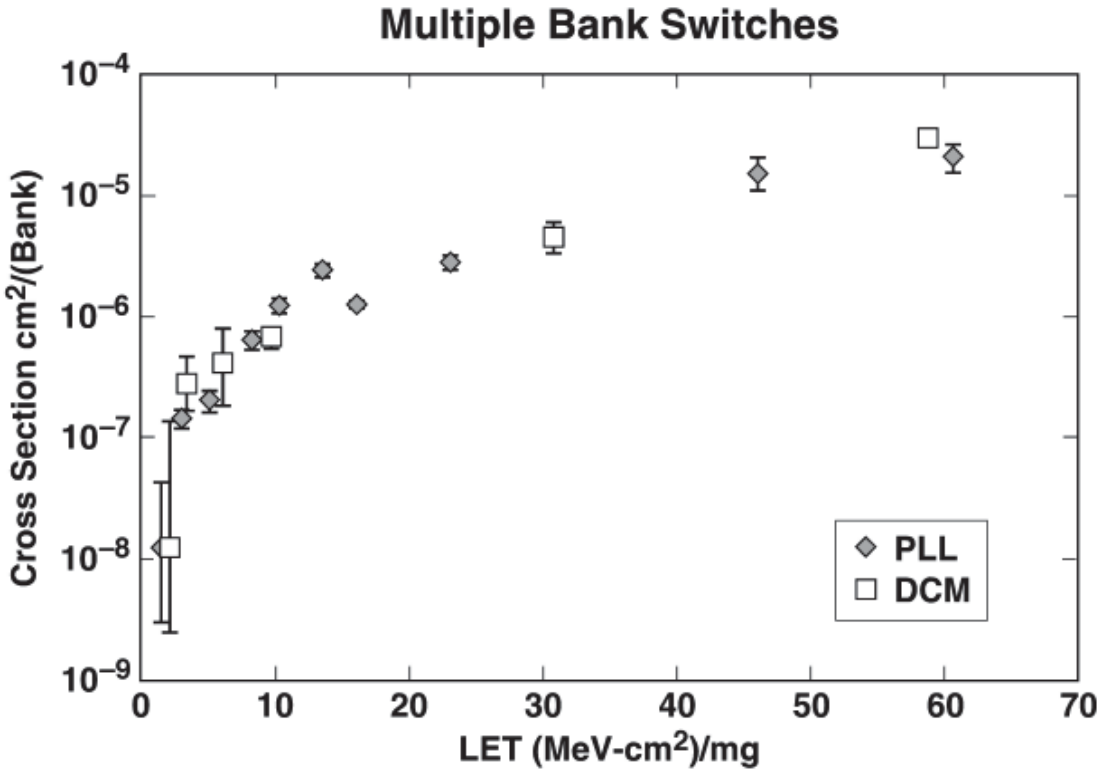

Figure 4.5-2. Cross-section vs. effective LET curve for multiple-bank switches is shown for events both on the PLL and DCM.

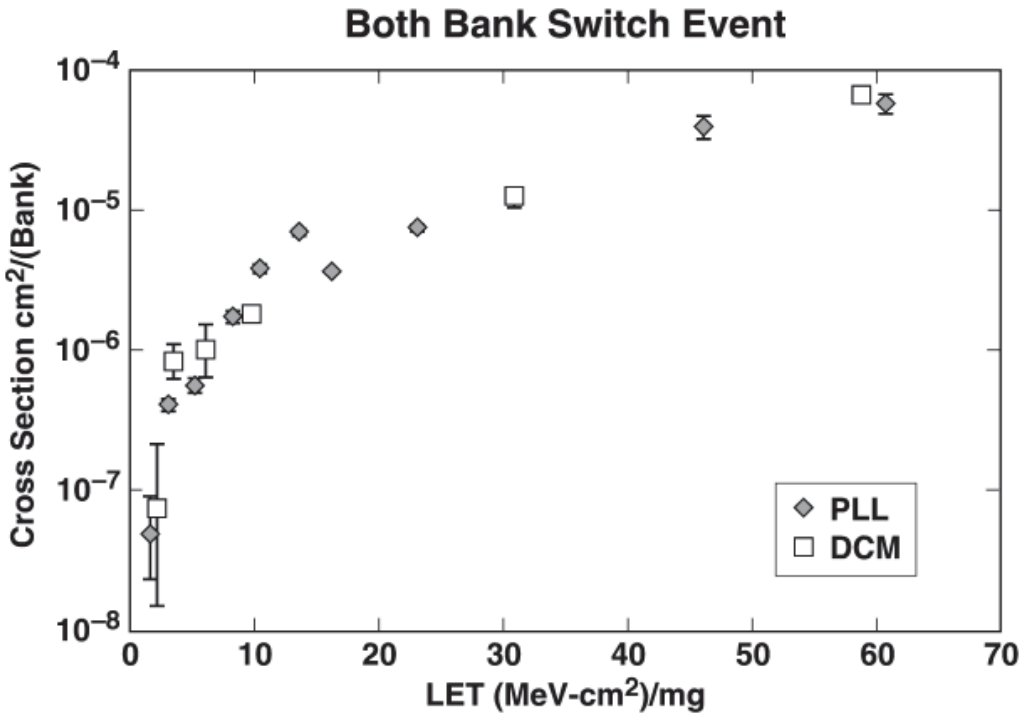

Figure 4.5-3. Cross-section vs. effective LET curve for both-bank switch events is shown for events both on the PLL and DCM.

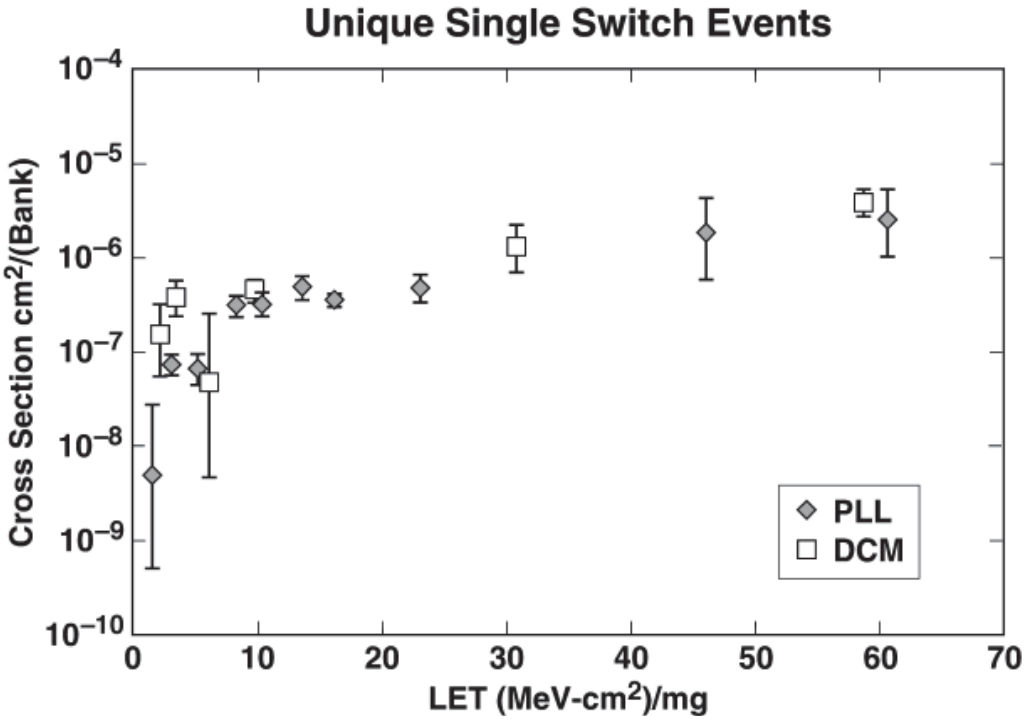

Figure 4.5-4. Cross-section vs. effective LET curve for unique single switch events is shown for events both on the PLL and DCM.

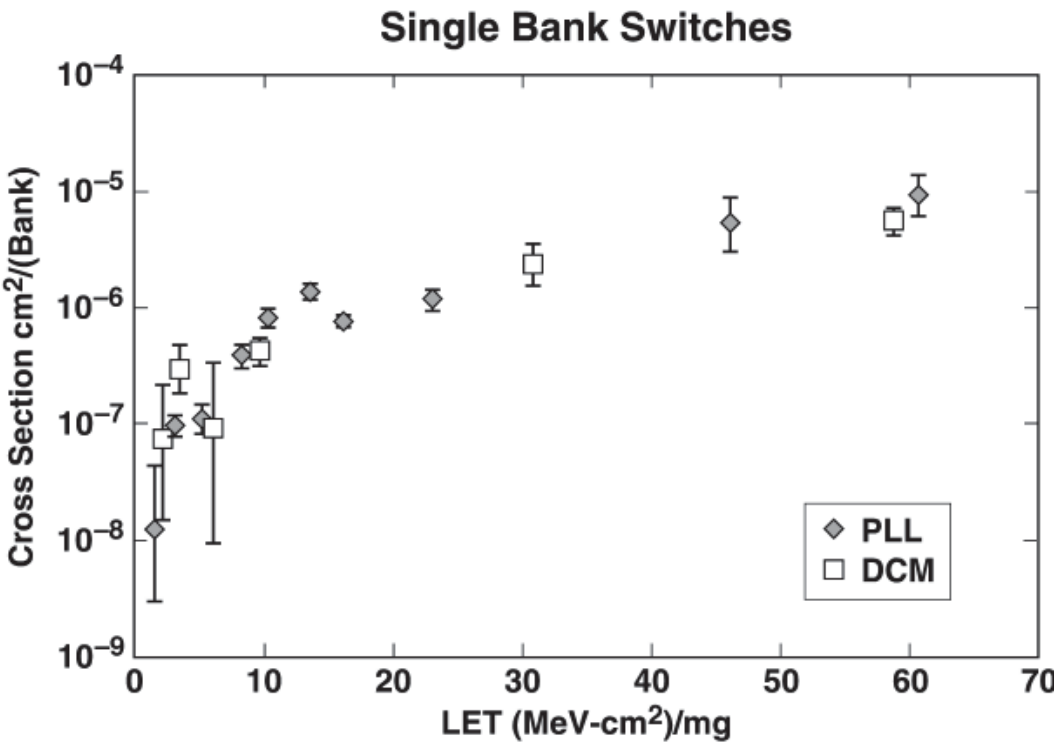

Figure 4.5-5. Cross-section vs. effective LET curve for single switch events is shown for events both on the PLL and DCM.

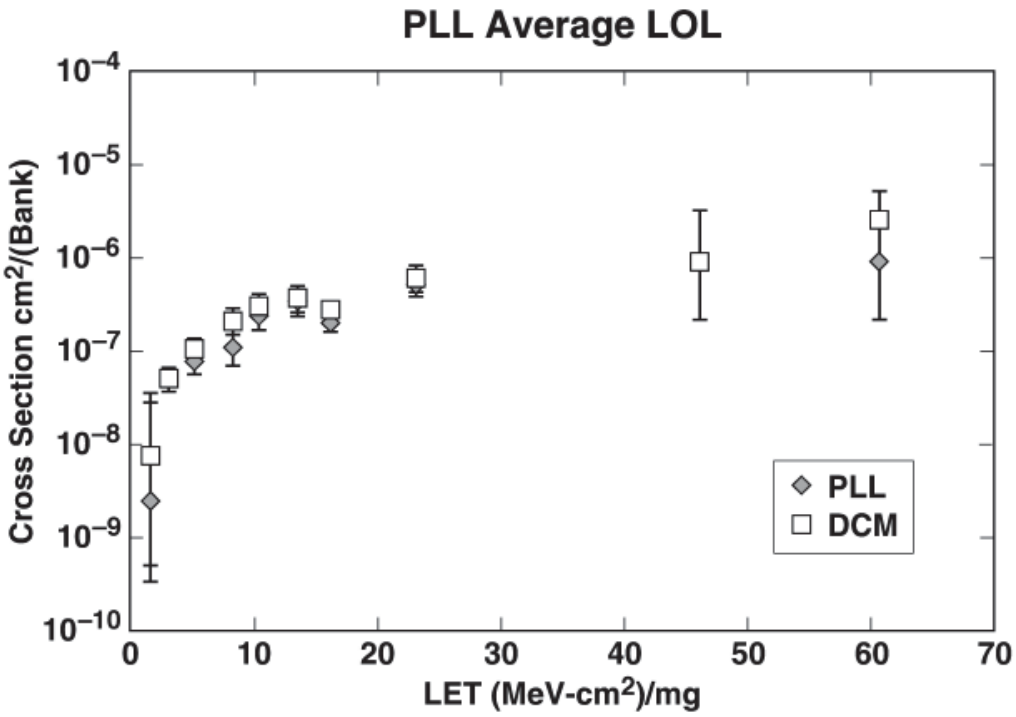

Figure 4.5-6. Cross-section vs. effective LET curve for loss of lock the PLL and DCM.

5.0 CONCLUSIONS AND FUTURE WORK

5.1 Conclusions

Two state-of-the-art commercial FPGAs, the Altera Stratix-IV and SiliconBlue iCE65 were characterized for SEL and found to be immune. The Xilinx Virtex-5 QV production device was also tested for SEL and, as expected, shown to be immune as well. Furthermore, the flip-flops and CMTs were extensively characterized. Between the characterization presented in last year's document [8], and the data shown here, many of the fundamental building blocks and IP have been characterized for SEE and shown to be relatively hardened compared to previous generations of SRAM-based reconfigurable FPGAs. Additionally, a great deal of effort has been put forth by other consortium members to characterize the DSPs and MGTs. While this data is not shown in this report, it is available in the literature.

5.2 Future Work

With the change in SEL response observed in the Altera devices, FY12 will target the full SEE evaluation of the Altera Stratix-V device. As for the Xilinx Virtex-5 QV devices, much of the characterization of individual components have been defined; however, how to fold the individual upset rates into a system error rate should be investigated. Furthermore, while the device is likely fairly immune to proton-induced SEE, a full characterization still needs to be performed.

6.0 REFERENCES

[1] Allen, G., "A Summary of Preliminary Xilinx Single Event Immune Reconfigurable Field Programmable Gate Array (SIRF) Single Event Effect (SEE) Testing," October 2009.

[2] Altera Corporation. Stratix-IV Device Family Overview. (January 2011). Retrieved from http://www.altera.com/literature/hb/stratix- iv/stx4_siv51001.pdf

[3] Document within private email from Sam Minger of Seakr Engineering. Sent July 1, 2005.

[4] Edmonds, L., "SEU Rate Estimates for Devices Exhibiting Dual-Node Susceptibility," JPL report, September 2007.

[5] ASTM, "Standard Guide for the Measurement of Single Event Phenomena (SEP) Induced by Heavy Ion Irradiation of Semiconductor Devices," ASTM Guideline F 1192-00, October 2000.

[6] Allen, G.R. and Swift, G.M. "Single Event Effects Test Results for Advanced Field Programmable Gate Arrays," 2006 IEEE Radiation Effects Data Workshop.

[7] Johnston, A.H. "The Influence of VLSI Technology Evolution on Radiation-Induced Latchup in Space Systems," IEEE Trans. Nucl. Sci., vol. 43, no. 2 (1996).

[8] Allen, G. "A Summary of Xilinx XQR5VFX130 Reconfigurable Field-Programmable Gate Array, Single-Event Effect (SEE) Testing." October 2010.

[9] Silicon Blue Technologies. iCE65 Ultra Low-Power Mobile FPGA Family Datasheet. (October 12, 2010). Retrieved from http://www.siliconbluetech.com/media/downloads/iCE65 Datasheet.pdf